High Performance Spaceflight Computing (HPSC)
an Avionics Formulation Task

Study Report

EXECUTIVE SUMMARY

Richard Doyle, Center Lead, *JPL*
Montgomery Goforth, Center Lead, *NASA JSC*
David Guibeau, Center Lead, *NASA KSC*
Michael Lowry, Center Lead, *NASA ARC*
Wesley Powell, Center Lead, *NASA GSFC*

Raphael Some, Technical Editor, *JPL*
Wesley Powell, Technical Editor, *NASA GSFC*

submitted by

Richard Doyle, Task Lead
Jet Propulsion Laboratory

submitted to

Garry Qualls, Program Investigator
Steve Gaddis, Program Manager
NASA Office of the Chief Technologist
Game Changing Development Program

October 22, 2012

The NASA HPSC Team

JPL:
Richard Doyle (Center Lead)
William Whitaker (Study Lead)
Nazeeh Aranki
Larry Bergman
Leonard Day
Kim Gostelow
Yutao He
Didier Keymeulen
David Rennels
Raphael Some (Technical Editor)
Carlos Villalpando

ARC:
Michael Lowry (Center Lead)
Matthew Scarpino

GSFC:
Wesley Powell (Center Lead, Technical Editor)
Thomas Flatley
Matt Holland
Michael Johnson
Semion Kizhner
Daniel Mandl
Jennifer Valdez
Jonathan Wilmot
Mark Wong

JSC:
Montgomery Goforth (Center Lead)

KSC:
David Guibeau (Center Lead)

1 Executive Summary

The goal of this formulation activity is to provide the information required by the Game Changing Development Program (GCDP) of the NASA Office of the Chief Technologist (OCT) Space Technology Program (STP) to develop a significantly improved spaceflight computing capability for NASA missions.

Space flight computing is a key resource in NASA space missions and a core determining factor of spacecraft capability, with ripple effects throughout the spacecraft, end-to-end system, and the mission; it is aptly viewed as a "technology multiplier" in that advances in onboard computing provide dramatic improvements in flight functions and capabilities across the NASA mission classes, and will enable new flight capabilities and mission scenarios, increasing science and exploration return per mission-dollar.

Space-qualified computing technology, however, has not advanced significantly in over ten years and the current state of the practice fails to meet the near- to mid-term needs of NASA missions. This is especially apparent in areas such as: Entry, Descent & Landing (EDL); Autonomous Rendezvous & Docking (AR&D); other forms of autonomous onboard Guidance, Navigation & Control (GN&C); autonomous mission planning; and science data processing and product generation for high data rate instruments.

In the past three years, a revolution has taken place in computing architectures and in semiconductor technology. Multi-core processing has enabled a two to three order-of-magnitude improvement in performance. Furthermore, state of the art semiconductor processes used in computer chip manufacturing exhibit a high degree of tolerance to two of the three broad categories of space radiation effects: Total Ionizing Dose (TID) and Single Event Latchup (SEL). The third type of radiation effect, Single Event Upset (SEU), can be mitigated by a set of techniques known as "Radiation Hardening by Design," or RHBD. The ability to develop radiation tolerant, space qualifiable processor chips inexpensively, compared to historical costs, by using relatively low-cost, commercially available semiconductor manufacturing processes and SEU mitigating RBHD circuits has been broadly recognized by the community, and several efforts are currently underway to develop new space qualifiable computers for future missions. NASA's unique requirements, however, driven by extreme power and reliability constraints, will not be met by these ongoing efforts as they are largely aimed at relatively benign environments with little to no consideration of fault tolerance, extreme long life, severe power and energy constraints, or deep space mission scenarios. This confluence of need and opportunity provide the impetus to develop a truly game changing technology for future NASA missions that will not be provided by any other means.

In formulating this study, as documented in the Formulation Activity Development Document (FADD), several key questions were defined:

- What are the paradigm shifting NASA space-based applications that drive flight computing?
- What are the requirements imposed on flight computing by these applications?
- What are the current or near-term offerings, and will they provide the required capabilities?
- Which computing architecture(s) provide the maximum return on investment for NASA?
- How can NASA most effectively invest its limited resources to effect the development of the required technology?

1.1 What are the paradigm shifting NASA space-based applications that drive flight computing?

To answer this question, a series of workshops was held with scientists and engineers from JSC, GSFC and JPL. Both SMD and HEOMD mission applications were addressed. The approach taken was to identify the high priority capabilities and mission scenarios, or "use cases" from future missions, and to focus on those use cases that required high performance computing. Table 1-1 lists the identified HEOMD and SMD use cases. A straightforward conclusion is that a high performance spaceflight computer will indeed be game changing because the capability is needed for many planned space missions across the agency, and will enable new and dramatic mission applications that are strongly desired by advanced mission planners.

Table 1-1 — Future NASA Mission Application Use Cases Requiring High Performance Spaceflight Computing

HEOMD Use Cases	SMD Use Cases
Cloud Services	Extreme Terrain Landing
Advanced Vehicle Health Management	Close Proximity Operations/Formation Flying
Crew Knowledge Augmentation	Fast Traverse
Improved Displays and Controls	New Surface Mobility Systems
Augmented Reality for Recognition and Cataloging	Imaging Spectrometers
Tele-Presence	Radar
Automated Guidance, Navigation and Control	Low Latency Products for Disaster Response
Human Movement Assist	Space Weather
Autonomous & Tele-Robotic Construction	Autonomous Mission Planning
	Immersive Environments - Science Ops/Outreach

1.2 What are the requirements imposed on the flight processor by these applications?

To answer this question, a series of workshops and discussions was held with engineers from JSC, GSFC, KSC and JPL, including the engineers and scientists who developed the use cases as shown in Table 1-1 above. By characterizing the required computing, the environment, the criticality of the application, and the system constraints, the computing system and processor chip requirements were derived for each use case application. In doing so, over 60 variations of the applications shown in the table above were examined. To reduce this to a manageable number, the most stressing and cross cutting of these applications were condensed into ten representative applications, which we termed '"Eigen-Applications". Each Eigen-Application represents a broad class of mission applications with similar requirements. Table 1-2 lists the Eigen-Applications and some of their salient requirements. The table indicates those applications requiring a high percentage of Digital Signal Processing (DSP), a high percentage of General Purpose Processing (GP), amenability of the application to parallelization (P), a requirement to run the application in a power constrained environment (LP), and whether the application is mission- or life-critical (MC).

Table 1-2 – Eigen-Application Summary Table.
Eigen-Applications Represent Classes of High Performance Applications. The table also shows the salient requirements for each application class.

Eigen-App	Throughput	DSP	GP	P	LP	MC
1	1-10 GOPS, 1-7 GFLOPS	X	X	X	X	
2	1-10 GOPS, 1-7 GFLOPS		X	X	X	X
3	10-50 GOPS, 7-35 GFLOPS	X	X	X	X	X
4	10-50 GOPS, 7-35 GFLOPS	X	X	X	X	
5	10-50 GOPS, 7-35 GFLOPS		X	X	X	X
6	10-50 GOPS, 7-35 GFLOPS		X	X	X	
7	50-100 GOPS, 35-70 GFLOPS	X	X	X	X	X
8	50-100 GOPS, 35-70 GFLOPS	X	X	X	X	
9	50-100 GOPS, 35-70 GFLOPS		X	X	X	X
10	50-100 GOPS, 35-70 GFLOPS		X	X	X	

To aid in the analysis of computing system capabilities, in accordance with the FADD, a set of Key Performance Parameters (KPPs) were defined. Acceptable levels of performance (i.e. thresholds) and stretch goals for these KPPs were defined in accordance with the values in Table 1.2. The Key Performance Parameters are:

- Computational Performance
- Fault Tolerance
- Radiation Tolerance
- Power and Energy Management
- Programmability and Flight Software Applicability
- Flight Software Verification and Validation (V&V)
- Interoperability
- Extensibility and Evolvability
- Cross-cutting potential across NASA missions
- Non-recurring cost
- Recurring cost

1.3 What are the current or near-term offerings, and will they provide the required capabilities?

To answer this question the HPSC team surveyed the currently available space computing platforms, as well as those currently in development and expected to be available within 3-4 years. The processors were grouped into categories by architectural type and the capabilities of these systems compared to the requirements. The result of this "gap analysis" was a determination that, while some future NASA mission requirements could be met by space computing systems currently in development, the large majority of interesting future mission scenarios would not be realizable without additional flight computing system development. This finding is not particularly surprising considering the extreme nature of NASA's missions compared to the commercial and military missions at which most space computing platforms are aimed. Table 1-3 lists the types of processors evaluated and the results of the gap analysis.

Table 1-3 – Gap Analysis.
No currently available or in-development processors meet requirements

Processor Category	Analysis Results
Rad Hard Single Core General Purpose Processors	Insufficient performance
Rad Hard Multicore Processors	Insufficient performance, high power, insufficient software development tools
Rad Hard SIMD Processors	High power, no GP capability – needs host processor
Rad Hard Reconfigurable Computers	Insufficient GP capability, difficult to program, high power
Redundant COTS Processors	Insufficient performance, high power, poor real time performance, excessive software complexity
Redundant COTS System Architectures	Insufficient performance, high power, mass, volume, excessive software complexity

1.4 Which computing architecture(s) provide the maximum return on investment for NASA?

To answer this question, a set of emerging mainstream computing architectures were identified. For each architecture, an extant exemplar of the state of the art was chosen in order to provide a specific instance of that architecture for use in the analysis. Utilizing the exemplar, each architecture was evaluated against each Eigen-Application, as well as the KPPs, by a team comprising engineers from GSFC, JPL, and ARC. The computing architectures identified for this study were:

- Rad Hard General Purpose Multicore
- COTS General Purpose Multicore
- DSP Multicore – COTS and Rad Hard
- Graphics Processor Unit (GPU) – COTS and Rad Hard
- Reconfigurable Computers – COTS and Rad Hard

Due to the prevalence and sharp focus in recent years on general purpose multicore architectures, in both the commercial and academic communities, COTS and Rad Hard General Purpose Multicore were treated in their own subsection. In all other cases, both COTS and radiation-hardened versions of the architecture are discussed in the same subsection.

The results of the individual architecture analyses were then weighted and traded to generate a final conclusion. The winner, by a significant margin, was the general purpose rad-hard multicore. Table 1-4 shows the results of the trade with weighted, un-weighted and scores above mean. In all cases, Rad Hard General Purpose Multicore scores significantly higher than the other architectures. Additional sensitivity analyses were performed using alternative weightings with similar results. As in the gap analysis explained above, this is not a surprising conclusion when viewed in the context of the full suite of NASA mission requirements.

Table 1-4 – Architecture Trade Studies.
In all cases, Rad Hard General Purpose Multicore has a significantly higher score and no disqualifying deficiencies.

Key Performance Parameter (KPP)	Rad-hardened General Purpose Multicore	Rad-hardened DSP Multicore	Rad-hardened Reconfigurable Computing	COTS-based Multicore	Rad-hardened Graphics Processing Units
Unweighted KPP Scores	52.1	37.8	40.9	39.3	38.0
Weighted KPP Scores	437	319	343	341	329
#KPP Scores Above Mean	12/12	4/12	7/12	6/12	5/12

In a broad sense, one can view the development of a radiation hardened general purpose multicore processor as “filling in the missing piece”. Some processors, such as reconfigurable FPGAs and DSP/SIMD machines, are becoming available in radiation hardened forms, and could be used in conjunction with a rad hard general purpose multicore processor to build “hybrid architectures”. Such machines may be useful for specific missions where there is a preponderance of the type of processing at which these architectures are aimed. For example, while a general purpose multicore machine can perform digital signal processing tasks reasonably well, it cannot match the performance of a dedicated DSP co-processor. The DSP co-processor, meanwhile, cannot stand alone. It requires a general-purpose host processor to manage its operation. The general purpose multicore can serve this function. Just a fraction of a general purpose multicore architecture’s processing capability might be dedicated to managing the DSP, and, if power constrained, most of the other GP cores might be turned off during DSP operation. Similarly, an FPGA, acting as a reconfigurable co-processor, can be managed by the general purpose multicore. Architectures like this are becoming more prevalent in the industry and may well result in standardized hybrid computer architectures in the near future.

1.5 How can NASA invest its resources to effect the development of the required technology?

To address this question, the team has drafted a proposed program plan. Inasmuch as a BAA type solicitation may be assumed, the team has also drafted a set of preliminary BAA requirements. The recommended program is divided into two portions, a competed portion and a directed portion. The competed portion acquires the processor hardware and the accompanying system software from industry via a BAA. The directed portion of the plan uses NASA application expertise to develop application-dependent “middleware,” software that logically resides between the operating system and the applications, and provides services to the application, which allow the application programmer to ignore the details of fault tolerance, power management and parallel program code and data distribution. The broad outline of the program plan is shown below.

- Competed Portion
  - — Solicit hardware from industry via a BAA
    - ♦ Option to seek innovative solutions
  - — Scope is the rad-hard general-purpose multi-core test board, including
    - ♦ Test board with bundled RTOS and FSW development environment
    - ♦ Improvements in fault tolerance and power dissipation at the hardware level

Phasing:

- — Phase 0: Multicore Architecture Innovation (9-12 months)
  - ♦ Evidence of additional measurable benefit to NASA and sufficient technical risk retirement
- — Phase 1: Multicore System Architecture Design (6 months)
  - ♦ Develop, design and validate next-generation avionics processor architecture
- — Phase 2: Processor Hardware Design (18 months)
  - ♦ Layout, integrated circuit fabrication, package and test target processor chip
- — Phase 3: System Integration and Validation (12 months)
  - ♦ Conduct radiation and thermal testing to validate models
  - ♦ ***Shared Product: Integrate chip components and software elements into test boards***

- Directed Portion
  - — NASA develops system software (middleware) for allocating/managing cores for fault tolerance and energy management purposes
    - ♦ Support for software-based fault tolerance methods
    - ♦ Thread monitoring for efficient load tracking and powering on/off of cores
    - ♦ Dynamic operation to trade performance, energy and fault tolerance
  - — Driven by knowledge of the NASA applications

The HPSC formulation study team wishes to thank NASA OCT STP and GCDP for the opportunity to perform this study. Special thanks to Harry Partridge and Steve Gaddis. Our team is confident that the recommendations presented here will allow NASA to exploit this unique juncture of technology and computer evolution to provide revolutionary new capabilities to future NASA missions.